

HAProxy

Powering Your Uptime

ALOHA Load-Balancer - Application Note

NFS high availability

Document version: v1.1

Last update: 4th March 2014

EMEA Headquarters

3, rue du petit robinson

ZAC des Metz

78350 Jouy-en-Josas

France

<http://www.haproxy.com/>

Purpose

Ensure high availability of a NFS service.

Limitation

This note does not explain how to synchronize data between both NFS servers.



In order to preserve file system integrity, usage of NFS servers will be in **Active/Passive mode only**

Complexity



Versions concerned

- Aloha 4.2.10 and above
- Aloha 5.5 and above

Changelog

Version	Description
1.1	<ul style="list-style-type: none">- HAProxy Tech. theme update- minor changes
1.0	Initial release

Before starting

Please read the documents listed below:

- Gateway mode explained:
<http://blog.exceliance.fr/2011/07/29/layer-4-load-balancing-direct-server-return-mode/>
- Server configuration to be compatible with gateway mode:
[appnotes_0053_server_configuration_for_layer4_dsr_mode_en.pdf](#)

Synopsis

Web servers use a NFS share to access data they deliver to clients. Data hosted by the NFS servers are accessed through the **ALOHA load-balancer**.

In order to avoid limiting the performance of the NFS service, we are going to use layer 4 load-balancing in **gateway** mode (also know as DSR: **Direct Server Return**).

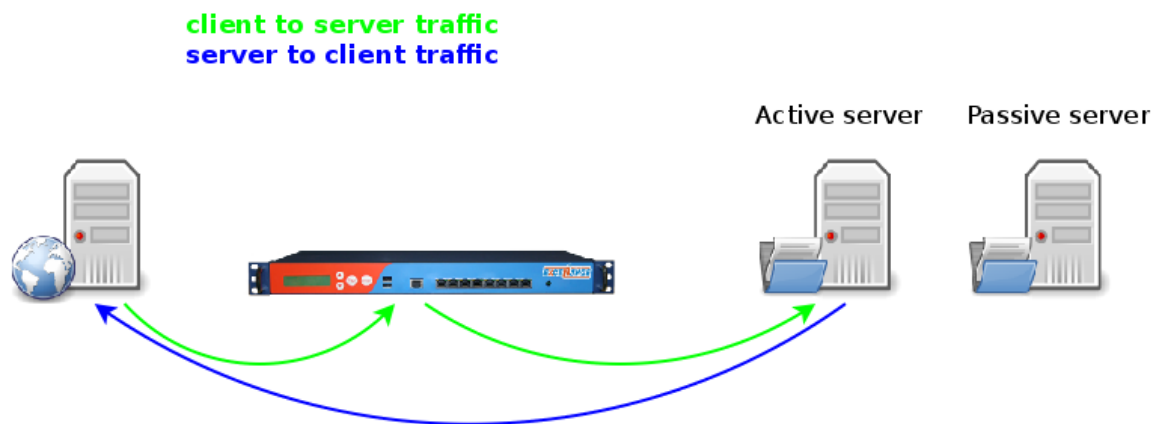
In this mode, the traffic back from the NFS server to the web server won't pass through the Aloha.



The gateway mode implies some modification on the **NFS server** side in order to prevent it from answering to ARP requests concerning the Virtual IP it hosts.

Diagram

The diagram below shows the flows for such architecture:



The client (here, the web server) pass through the **ALOHA** to access the NFS service. The NFS server talk directly to the client, bypassing the ALOHA load-balancer.

Configuration

NFS service

Since NFS can use random ports, the configuration must be done in three steps:

- Matching and routing NFS network flows
- Load-Balancing
- LVS service tuning to speed up the fail over



In order to make the service cleaner, it is possible to fix ports used by the NFS services on the server and update accordingly the load-balancing configuration in the **ALOHA**

Flow manager

- Click on the GUI **Flow** tab
- Add the lines below:

```
flow nfs director nfs
match iface eth0 dst 192.168.10.50
```

Layer 4 load-balancing

- Click on the GUI **LB layer 4** tab
- Add the lines below:

```
director nfs
balance roundrobin
mode gateway
check interval 10 port 2049 timeout 2
option tcpcheck
server nfs-01 192.168.10.100:2049 weight 10 check
server nfs-02 192.168.10.101:2049 sorry
```

LVS service tuning

When a fail over occurs, the **ALOHA** redirects automatically **new connections only**.

Established connections keep on being redirected to the server which managed them (and which is obviously currently unavailable).

Any NFS client waits up to 15 minutes before opening a new connection, which means the web server won't be able to deliver any data during that amount of time.

In order to speed up convergence, it is possible to configure the **ALOHA** to send a TCP RST packet to the client. That way, the client opens a new connection. To do this, we just need to enable two sysctls in LVS service configuration.

- Click on the GUI **Services** tab
- Click on the LVS service Edit icon and add the lines below:

```
sysctl expire_nodest_conn=1
sysctl expire_quiescent_template=1
```

- Save and restart the LVS service